

Depth Restoration of Hand-Held Transparent Objects for Human-to-Robot Handover

Ran Yu^{1*}, Haixin Yu^{1*}, Shoujie Li^{1*}, Yan Huang¹, Ziwu Song¹, Wenbo Ding^{1,2†}

Abstract—Transparent objects are common in daily life, while their optical properties pose challenges for RGB-D cameras to capture accurate depth information. This issue is further amplified when these objects are hand-held, as hand occlusions further complicate depth estimation. For assistant robots, however, accurately perceiving hand-held transparent objects is critical to effective human-robot interaction. This paper presents a Hand-Aware Depth Restoration (HADR) method based on creating an implicit neural representation from a single RGB-D image. The proposed method utilizes hand posture as an important guidance to leverage semantic and geometric information of hand-object interaction. To train and evaluate the proposed method, we create a high-fidelity synthetic dataset named TransHand-14K with a real-to-sim data generation scheme. Experiments show that our method has better performance and generalization ability compared with existing methods. We further develop a real-world human-to-robot handover system based on HADR, demonstrating its potential in human-robot interaction applications.

I. INTRODUCTION

RGB-D cameras have become popular for robotic perception due to their ability to capture 3D information, which is crucial for robot grasping [1] and manipulation [2]. However, transparent objects pose challenges in capturing correct depth information from RGB-D cameras due to light refraction and reflection [3]. This issue is particularly critical for service robots [4], which are deployed in scenarios where transparent objects widely exist. Given the growing demand for human-robot interaction, it's essential for robots to accurately perceive transparent objects held by humans and manipulate them safely. This further increases the need for precise and robust depth estimation of transparent objects.

Many attempts have been made to address the problem of transparent object depth estimation by introducing depth restoration models [3], [5], [6]. The general idea behind those methods is extracting information from the RGB and corrupted depth images to predict the correct depth value of

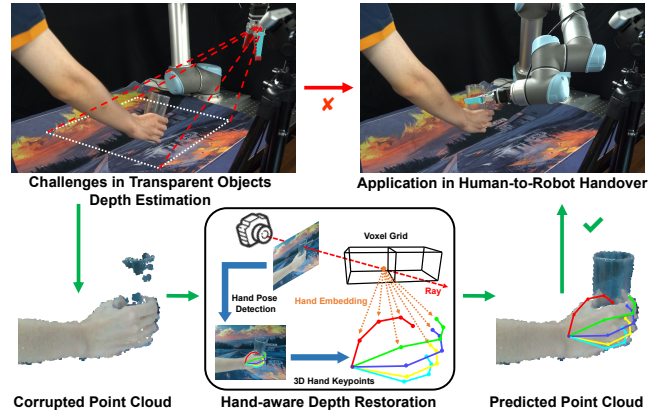


Fig. 1. RGB-D sensors meet challenges in estimating the depth of transparent objects, especially hand-held objects. We present a hand-aware depth restoration method that reconstructs the corrupted point cloud with the guidance of hand pose information. Our method can be used for human-to-robot handover, highlighting its application value.

transparent objects. However, restoring the hand-held transparent objects presents many new challenges and a research gap still exists. **First**, generating data for model training could be much harder in both real-world and simulation settings. At present, there exists no dataset supporting this area of research. **Second**, hand occlusions bring difficulties in depth restoration. Since objects' appearances are highly related to the hand geometrically, the same object may look different because of hand occlusions, which leaves a higher demand for model generalization ability. **Additionally**, previous studies on human-to-robot (H2R) handover [7], [8] claim their methods are struggled with transparent objects. It remains uncertain whether incorporating depth restoration can improve the overall performance of the transparent object handover.

In this paper, we aim to solve the aforementioned problems from the perspectives of data, algorithm, and application. A synthetic dataset named TransHand-14K is first proposed, which is used for hand-held transparent object perception research. We implement a real-to-sim data generation scheme to ensure the quality and fidelity of data. Based on the dataset TransHand-14K, we present a **Hand-Aware Depth Restoration (HADR)** method for hand-held transparent objects. This method incorporates hand-holding postures to facilitate object reconstruction. Results demonstrate that our method surpasses existing methods, especially in terms of object-level generalization ability. Furthermore, we introduce a human-to-robot handover workflow for transparent objects based on the proposed depth restoration method. This workflow is adaptable to a wide range of commonly encountered

*Authors with equal contributions. †Corresponding author.

¹Shenzhen Ubiquitous Data Enabling Key Lab, Shenzhen International Graduate School, Tsinghua University, Shenzhen 518055, China. (Emails: yur23@mails.tsinghua.edu.cn, yuhx21@tsinghua.org.cn, {lsj20, Huang-y24, song-zw20}@mails.tsinghua.edu.cn and ding.wenbo@sz.tsinghua.edu.cn)

²RISC-V International Open Source Laboratory, Shenzhen 518055, China.

This work was supported by Shenzhen Key Laboratory of Ubiquitous Data Enabling (No. ZDSYS20220527171406015), Shenzhen Science and Technology Program (JCYJ20220530143013030), Guangdong Innovative and Entrepreneurial Research Team Program (2021ZT09L197), Tsinghua Shenzhen International Graduate School-Shenzhen Pengrui Young Faculty Program of Shenzhen Pengrui Foundation (No. SZPR2023005), Shenzhen Higher Education Stable Support Program (WDZC20231129093657002), and Meituan.

Project website: <https://marcyu0303.github.io/HADR.github.io/>.

transparent objects and is capable of dynamically responding to human movements during the handover process.

In summary, the key contribution of our work includes the following three aspects: 1) a synthetic dataset TransHand-14K for transparent object perception research, 2) a novel depth restoration method HADR for hand-held transparent objects, 3) a human-to-robot transparent objects handover workflow deployed in a real-world system.

II. RELATED WORK

A. Depth Restoration of Transparent Objects

To minimize errors in the depth estimation of transparent objects using 3D sensors, various depth restoration methods are proposed, which utilize RGB and corrupted depth data to reconstruct the objects' shape. Some early works restore the depth with multiple camera viewpoints based on visual hull [9], stereo-view matching [10], and NeRF [11], [12]. While multi-view methods suffer from relatively long inference time, other works reconstruct transparent objects only using single-view. A global optimization method is introduced in [3] that restores corrupted depth by predicting objects' surface normals and edges. In [13], a local implicit neural representation method is proposed. Authors in [14] design an end-to-end depth completion network with Unet architecture called DFNet. To better extract features from input data, transformer architecture [15] is applied to some recent works such as TODE-trans [5] and SwinDRNet [6]. However, all aforementioned methods are designed to reconstruct transparent objects on the tabletop. Depth restoration of hand-held transparent objects remains an open problem, which is crucial for human-robot interaction.

B. Datasets for Transparent Objects Perception

Training a depth restoration model for transparent objects requires both corrupted and perfect depth data from RGB-D cameras. These data can be generated from either the real world or the simulation environments. A large-scale dataset is proposed in [3], which contains both synthetic data rendering with Blender and real-world data by spraying opaque material on transparent objects. In the work of [13], a synthetic dataset of transparent objects is generated by using the NVIDIA Omniverse platform [16]. Another large-scale real-world dataset called TransCG is introduced in [14], which generates data with an auto-collection pipeline. However, no dataset is currently available for grasped or hand-held transparent objects, as creating such a real-world dataset could require significant time and financial resources. Creating a synthetic dataset also presents challenges, particularly in generating realistic hand poses for objects with various shapes, which is crucial for narrowing the sim-to-real gap.

III. METHODOLOGY

A. TransHand-14K Dataset Generation

In this work, we propose a synthetic dataset specifically designed for hand-held transparent objects, named TransHand-14K. Generating such a dataset is more challenging than the conventional object-only datasets [3], [13].




Fig. 2. TransHand-14K data generation method. (a) A real-to-sim hand pose generation scheme is introduced. (b) We use a MANUS VR glove to capture the pose data. (c) Eight transparent objects are included in TransHand-14K.

Beyond considerations of background setting and lighting, it is also essential to account for the various handhold postures and their realism. Existing works of hand pose generation have some limitations [17]. For instance, the hand-held object dataset proposed by [18] only provides six types of grasping schemes for different objects. Additionally, the generated data suffers from problems of unrealistic background and hand poses.

To solve those problems, we propose a real-to-sim data generation scheme. As shown in Fig. 2(a), we first utilize a MANUS VR glove [19] (Prime X Haptic) to capture the hand pose data of holding transparent objects, and then map the twin hand poses in Blender [20]. The synthetic hand in Blender is built on a skeleton model as shown in Fig. 2(b). To reduce the difficulty of modeling in Blender, we adopt a 22-point skeleton model by creating a virtual point (number 2), which is different from the commonly used 21-point skeleton model. The skeletal control can be achieved by feeding joint angle data captured by the MANUS VR glove. After generating the handhold poses, we manually place the corresponding transparent object into the proper position. 3D scanning is used to model eight common transparent objects with different shapes into Blender for rendering. In addition to ensuring the realism of generated hand posture, the proposed real-to-sim data generation method can easily obtain various grasping poses of the same object, significantly reducing the difficulty of dataset construction.

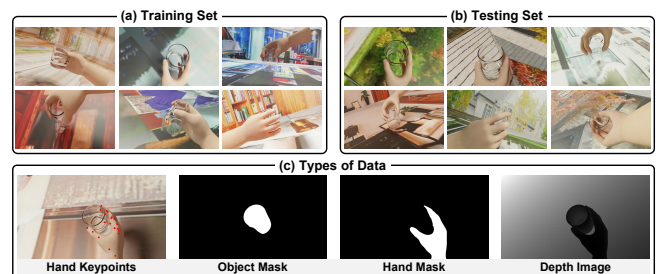


Fig. 3. Visualization of proposed dataset TransHand-14K.

For the scene and background setting, we follow the work of Trans6D-32K [21]. Our dataset contains 14100 images

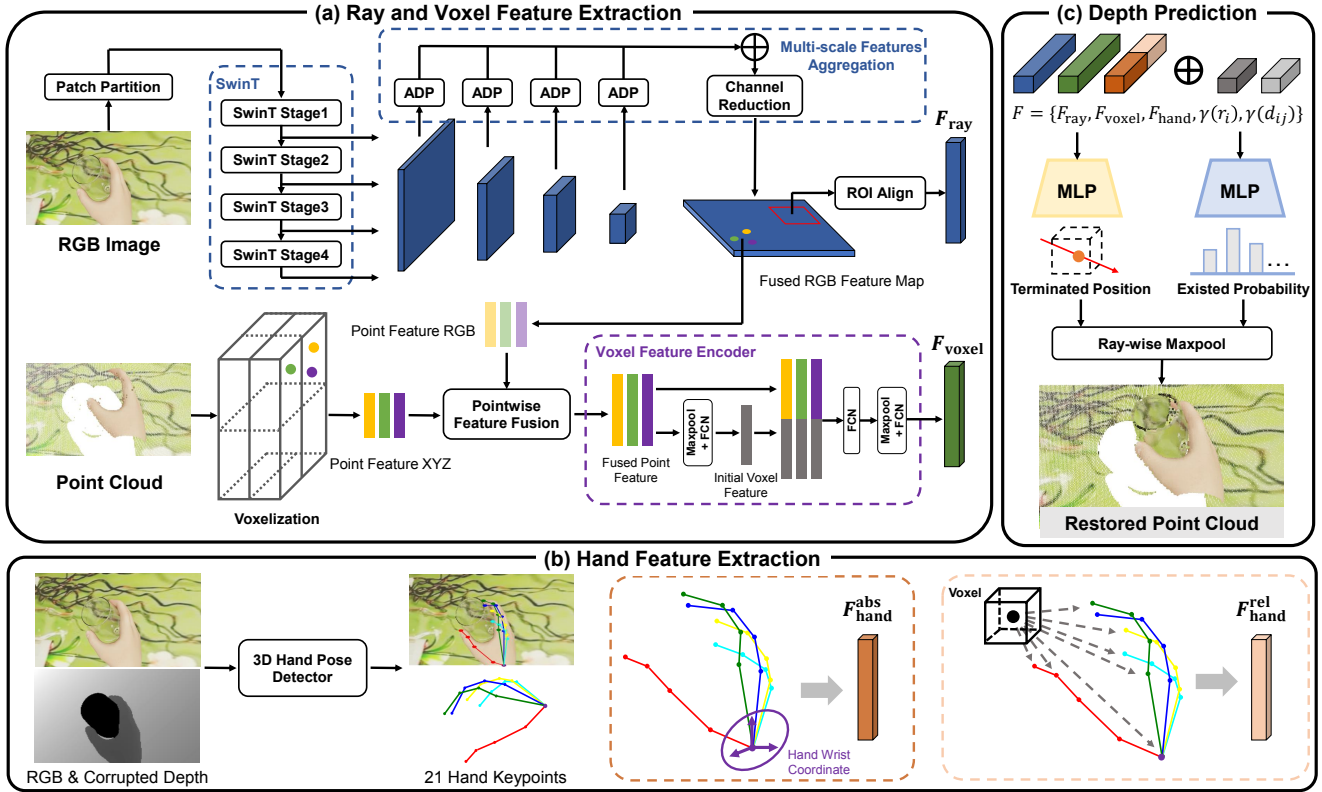


Fig. 4. Overview of HADR. (a) Ray and voxel features are generated from the RGB and corrupted point cloud. (b) We introduce the handhold pose as an important guidance for depth restoration. This feature is related to geometric and semantic information of hand-object interaction. (c) The terminated probability and position of each ray-voxel pair are predicted. The final restored depth is obtained by ray-wise maxpooling.

with a resolution of 1280×720 . As shown in Fig. 3, TransHand-14K contains RGB image, perfect depth image, objects mask, hand mask, and labels of the hand keypoints. In addition to depth completion, this dataset can be used for other transparent object perception tasks such as detection, recognition, segmentation, and 6D pose estimation.

B. Hand-Aware Depth Restoration

Ray-Voxel Implicit Neural Representation. We propose a depth restoration method named HADR for hand-held transparent objects. Inspired by the work of [13], our method is based on generating a local implicit neural representation of ray-voxel pairs. The corrupted depth image is first transformed into the point cloud and placed into a voxel grid space in the camera coordinate. A pixel I_i in the image can be viewed as a camera ray r_i which passes through a set of occupied voxels \mathbf{V}^{occ} . The ray-voxel pair [13] is defined as $\phi_{ij} = (r_i, v_j)$, where r_i is the camera ray direction for i -th pixel and v_j is the intersected voxel with index j .

Based on the definition of ray-voxel pair, restoring a missing depth value D_i is equivalent to finding out it exists in which voxel and the offset distance in this voxel, which can be expressed as

$$d_i = d_{ij'}^{in} + \delta_{ij'} r_i, \quad j' = \underset{j: v_j^{occ} \in \mathbf{V}_{r_i}^{occ}}{\operatorname{argmax}} p_{ij}^{end}, \quad (1)$$

where d_i is the predicted depth of ray r_i ; $d_{ij'}^{in}$ is the entering position of ray-voxel intersected point; $\delta_{ij'}$ is the offset distance; p_{ij}^{end} is the probability of ray r_i terminated in the

voxel v_j . This expression indicates that we can use neural networks to generate an implicit function for predicting the terminated position and probability of each ray-voxel pair, thereby restoring the accurate depth of transparent objects.

Ray and Voxel Feature Extraction. To construct the implicit neural representation, we first generate ray-wise and voxel-wise features for each ray-voxel pair. Fig. 4(a) shows the model architecture of extracting these features from the input RGB image and corrupted point cloud.

For the input RGB image I_C , we use the swin transformer [22] as the backbone to extract multi-scale image features $\{F_{rgb}^i\}_{i=1,2,3,4}$, where $F_{rgb}^i \in \mathbb{R}^{\frac{H}{4^i} \times \frac{W}{4^i} \times C}$ and $C = 96$ is the embedding dimension of first layer output. We then introduce a features aggregation module to fuse multi-scale features into a single dense color feature map $F_{rgb}^{dense} \in \mathbb{R}^{H \times W \times C_{rgb}}$, where $C_{rgb} = 32$. This is realized by first aligning different features into the same dimension of $\mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$, and concatenating them together. We then reduce the channels into C_{rgb} using a 1×1 convolution layer and conduct bilinear interpolation to resize the feature into $H \times W$. For the ray r_i , its corresponding ray feature can be extracted from F_{rgb}^{dense} using ROI alignment [23], with an input size of 8×8 and an output size of 2×2 . The ray feature F_{ray} is obtained by flattening the output of the ROI alignment.

For the voxel feature, early work [13] generates it by sending the position and color of the valid point cloud into a voxel-based PointNet [24], [25]. However, this approach

may not fully utilize the semantic information provided by the RGB image. To overcome this limitation, we introduce an early-stage point fusion method. We first encode the geometry embedding of point cloud $P_{xyz} \in \mathbb{R}^{N \times 16}$ using the relative position of points to their voxel centers. We then project the point cloud with camera intrinsic parameters onto the generated dense color feature map F_{rgb}^{dense} and generate point-wise color embedding $P_{rgb} \in \mathbb{R}^{N \times 16}$ through channel reduction. These two features are concatenated into the fused point feature P_{fuse} . The final voxel feature F_{voxel} is generated from P_{fuse} using a voxel feature encoder, which contains a set of voxel-wise maxpooling layers and fully connected layers [26].

Hand Feature Extraction. In addition to the ray and voxel feature, HADR incorporates hand pose as important guidance for depth restoration. One of our insights is that the hand pose can provide both geometric and semantic context, enhancing the accuracy and generalization ability of object reconstruction. As shown in Fig. 4(b), we extract the 21 hand keypoints with a pre-trained hand pose detector. In our experiment, a ResNet-18 [27] is adopted as the backbone to extract the hand feature map from the corrupted depth image, which is then transferred into hand keypoints $X_{hand} \in \mathbb{R}^{21 \times 3}$ through the weighted average regression [28].

Two types of hand features are derived from the detected hand keypoints. We first transfer the keypoints into the hand-wrist coordinate and obtain the absolute hand feature F_{hand}^{abs} . This feature represents the absolute hand-holding pose, which is invariant to the change of camera position and orientation. The second type of hand feature F_{hand}^{rel} is the relative position between each hand keypoint and the center of intersected voxel v_j . This feature represents the spatial relationship of hand-object interaction. Absolute and relative hand features are flattened and concatenated into the final hand feature F_{hand} .

Depth Prediction. As shown in Fig. 4(c), the final embedding feature can be formulated as Equation 2:

$$F = F_{ray} \oplus F_{voxel} \oplus F_{hand} \oplus \gamma(r_i) \oplus \gamma(d_{ij}), \quad (2)$$

where \oplus represents the feature-wise concatenation operator, $\gamma(r_i)$ and $\gamma(d_{ij})$ are the ray direction and voxel position embedding, which has shown to be important for the final accuracy [13]. Two MLPs are used to decode the feature into the predicted terminated position and the existing probability of each ray-voxel pair. We obtain the final predicted depth value by using ray-wise maxpooling [13], as expressed in Equation 1.

Training Objective. The loss function defined in Equation 3 is optimized for training:

$$L = w_{depth}L_{depth} + w_{prob}L_{prob} + w_{norm}L_{norm}, \quad (3)$$

where L_{depth} is the first-order norm between ground truth and predicted depth value; L_{prob} is the cross entropy loss between the ground truth and predicted voxel existing probability; L_{norm} is the loss of the surface normal measured by the cosine similarity; w_{depth} , w_{prob} , and w_{norm} are the weights of different loss components.

C. H2R Handover Workflow of Transparent Objects




Fig. 5. Overview of the proposed handover workflow. The whole handover process is divided into three stages: 1) Wait & Observe, 2) Approach & React, and 3) Grasp & Retrieve.

Based on the proposed hand-aware depth restoration method, we further develop a handover workflow for transparent objects. As shown in Fig. 5, we divide the whole H2R handover process into three stages: 1) Wait & Observe, 2) Approach & React, and 3) Grasp & Retrieve.

Wait & Observe. This is the starting phase of the whole handover task. The robot is waiting at the home position, and the depth restoration module is activated to restore the corrupted depth. GraspNet [1] is used to generate a set of 6-DoF grasps G given the restored point cloud of the transparent object. Each potential grasp g is associated with a score s . We first select the grasps on the transparent objects based on the generated object mask and then filter the grasps based on the collision detector. For the remaining available grasp, we rescore them based on the metric in Equation 4:

$$C = w_s s + w_{r2h} \mathbf{v}_{r2h}^\top (R_g \mathbf{v}_{r2h}) + w_{u2d} \mathbf{v}_{u2d}^\top (R_g \mathbf{v}_{u2d}), \quad (4)$$

where R_g is the rotation of grasp pose, \mathbf{v}_{r2h} denotes the unit vector pointing from the robot to the human, and \mathbf{v}_{u2d} denotes the unit vector pointing from the up to the down. We design the metric in this way because a grasp facing toward the user and vertically pointing downward is preferred. w_s , w_{r2h} , and w_{u2d} are the weights, which are set to be $w_s = 1$, $w_{r2h} = 0.4$, and $w_{u2d} = 0.2$. The grasp pose with the highest score will be chosen, and the corresponding pre-grasp pose is defined as 10cm back from the predicted grasp along its z-axis.

Approach & React. In this stage, the goal of the robot is to approach the pre-grasp pose. Since the human user may move their hand during this stage, the robot should react to the movement in a closed-loop manner. We assume that the absolute holding pose will not be changed significantly

during the handover and realize the real-time grasp pose adjustment by tracking the hand keypoints.

We denote the initial and current hand keypoints as X_i and X_c in the dimension of $\mathbb{R}^{21 \times 3}$. The translation vector is computed as $T = X_c^0 - X_i^0$, where X_i^0 and X_c^0 represent the initial and current wrist position. For the rotation, we perform the Singular Value Decomposition on the covariance matrix H , defined in Equation 5,

$$H = (X_i - X_i^0)^\top (X_c - X_c^0) = U\Sigma V^\top. \quad (5)$$

The rotation matrix is then calculated as $R = VU^\top$, where R should be adjusted by flipping the sign of the last column if its determinant is negative. The translation vector and rotation matrix are then applied to the initial pre-grasp pose to obtain the current one.

Grasp & Retrieve. Once the robot reaches the pre-grasp position, it will grasp the object in an open-loop manner. We move the final grasp position 5cm forward along the z-axis as we assume the human user will adapt to the robot’s grasp motion. Once the gripper closes successfully, the robot will retrieve and drop the object.

IV. EXPERIMENTS

A. Depth Restoration Experiments on TransHand-14K

Evaluation Metrics. We adopt four commonly used metrics for depth estimation [3]. 1) **RMSE**: the root mean square error between the predicted result and ground truth. 2) **REL**: the mean value of absolute relative difference. 3) **MAE**: the mean absolute error between the predicted result and ground truth. 4) **Threshold δ** : the percentage of data smaller than the given threshold δ . $\max(d_i/d_i^*, d_i^*/d_i) < \delta$, where d represents the predicted depth, d^* represents the ground truth depth, and $\delta \in \{1.05, 1.10, 1.25\}$. All metrics are evaluated on the transparent objects area.

Implementation Details. We set the image resolution to be 224×224 for both RGB and depth input. The RGB input is further augmented by adjusting the color space and incorporating additional elements such as bright patches, blur, and noise. We train the network for 100 epochs and choose Adam as the optimizer with a learning rate of $1e-3$ for the first 80 epochs and $1e-4$ for the last 20 epochs. The number of voxels used for restoration is 8^3 . We set $w_{\text{depth}} = 200$, $w_{\text{prob}} = 10$, and $w_{\text{norm}} = 0.5$ in training.

Comparison with SOTA. We compare the proposed HADR with several state-of-the-art methods, including TODE-Trans [5], TransCG [14], SwinDRNet [6], and LIDF [13]. All baselines are trained on the proposed dataset TransHand-14K. We split the training, validation, and testing set according to the ratio of 7:2:1.

Table I reports the comparison results. We first compare methods on the known category objects, where models are trained and evaluated on all eight different shapes of transparent objects. Our method achieves the best performance in terms of RMSE, MAE, $\delta_{1.05}$, and $\delta_{1.10}$. For the other two metrics, our method is slightly weaker than SwinDRNet. For the unknown category evaluation, we choose six objects for training and the other two for evaluation. Our method




Fig. 6. The setup of our real-world robot experiment.

significantly outperforms other methods in terms of all metrics, demonstrating the strong generalization ability of the model. This better performance can be attributed to our hand-aware design, which enables the model to effectively transfer similar hand-object interaction patterns to unknown objects. Additionally, unlike LIDF [13], our method can directly produce high-quality final results without needing a depth refinement process.

TABLE I
QUANTITATIVE COMPARISON WITH STATE-OF-THE-ART METHODS.

Methods	RMSE↓	REL↓	MAE↓	$\delta_{1.05} \uparrow$	$\delta_{1.10} \uparrow$	$\delta_{1.25} \uparrow$
Known Category Evaluation						
TODE-Trans [5]	0.024	0.056	0.017	57.03	82.55	98.61
TransCG [14]	0.012	0.026	0.008	86.91	96.96	99.76
SwinDRNet [6]	0.009	0.015	0.005	94.82	98.26	99.87
LIDF+Refine [13]	0.014	0.029	0.010	85.18	96.08	99.60
HADR (Ours)	0.009	0.016	0.005	95.22	98.44	99.81
Unknown Category Evaluation						
TODE-Trans [5]	0.052	0.055	0.037	62.86	89.26	97.21
TransCG [14]	0.027	0.064	0.020	51.54	78.54	98.79
SwinDRNet [6]	0.022	0.049	0.015	65.07	86.79	98.79
LIDF+Refine [13]	0.026	0.063	0.021	52.66	78.69	97.90
HADR (Ours)	0.018	0.042	0.014	72.74	90.06	99.05

TABLE II
ABLATION STUDY OF FEATURES AND MODULES.

	RMSE↓	REL↓	MAE↓	$\delta_{1.05} \uparrow$	$\delta_{1.10} \uparrow$	$\delta_{1.25} \uparrow$
w/o Hand Feature	0.012	0.019	0.006	93.35	97.79	99.70
w/o MS Aggregation	0.011	0.020	0.007	92.25	97.49	99.70
w/o Point Fusion	0.009	0.017	0.006	94.40	98.37	99.80
2D Hand Feature	0.010	0.017	0.006	94.09	98.30	99.77
Full Model	0.009	0.016	0.005	95.22	98.44	99.81

Ablation Study. We conduct ablation experiments to analyze the different components of the proposed method. To assess the impact of hand-ware on performance, we compare the model with two other versions: the model excludes the hand-pose feature and uses 2D hand keypoints as the feature. As shown in Table II, the hand pose feature contributes to the depth restoration result. Although the 2D keypoints feature produces a slightly inferior performance compared with the 3D version, in real-world applications, obtaining the 2D keypoints could be relatively easy and robust in some situations. We also evaluate the effectiveness of multi-scale feature aggregation and point-wise feature fusion. Results indicate that removing these two modules will lead to a drop in overall performance.

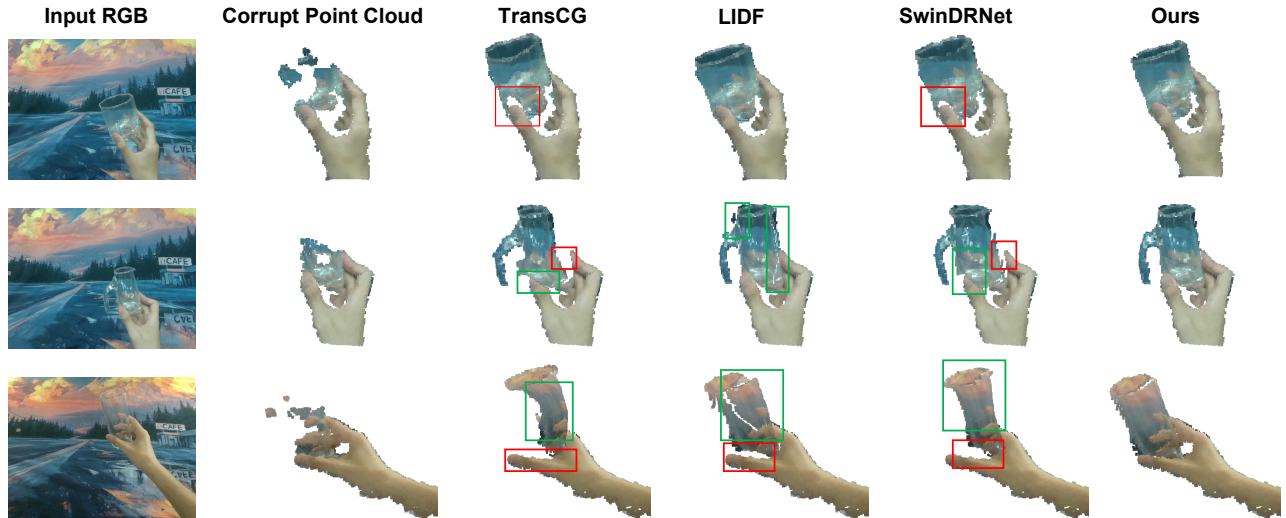


Fig. 7. Qualitative evaluation in real scenarios. We compare the proposed HADR to state-of-the-art methods. Defects including deformation (in the color of green) and spatial shifting (in the color of red) are highlighted. For better visualization, we remove point cloud of the background.

B. Real-World Experiments

We conduct real-world experiments to verify the effectiveness of HADR and handover workflow.

Experiment Setup. Fig. 6 illustrates the setting of our real robot system. A UR5 robot arm equipped with a soft gripper is rigidly mounted on the table, which is used to perform the robot handover task. For perception, we employ a RealSense D435 camera with an RGB-D stream. The camera is placed in a third-person view with extrinsic calibrated. A human participant stands on the opposite side of the table, holding a transparent object in hand and passing it to the robot. We choose the tablecloth with intricate patterns and rich colors as the background to examine the robustness of our method.

Implementation. In addition to implementing the proposed depth restoration method, some other perception models are used in the real-world experiment. To obtain the segmented area of transparent objects, we train a segmentation model in the Unet architecture with a ResNet-34 [27] as the encoder using the TransHand-14K dataset. For the hand pose detector, we find the weighted average regression from the depth image is vulnerable to occlusions and noise. Instead, we first generate 2.5D hand keypoints using an off-the-shelf RGB-based hand pose detector [29], and then extend it into 3D keypoints using the depth value of the wrist keypoint position. The pre-trained GraspNet-baseline [30] is used to generate 6-DoF grasping poses. All the aforementioned models are implemented using a computer with an NVIDIA 3080 GPU. In the experiment, the whole perception pipeline achieves an inference speed of 11 FPS, allowing real-time depth restoration.

Qualitative Evaluation in Real Scenarios. To better illustrate the sim-to-real performance of the proposed method, a qualitative evaluation is conducted by comparing it to the state-of-the-art methods in real scenarios. As shown in Fig. 7, it can be observed that the reconstructed object point cloud exists deformation (highlighted with green boxes) and spatial shifting (highlighted with red boxes). Results demonstrate that our method reconstructs transparent objects with greater

accuracy and robustness in real-world settings.

Handover Experiment. We conduct a human-to-robot handover experiment using the proposed handover workflow. In the experiment, eight transparent objects with different shapes are used. To better illustrate the generalization ability of our model, three of them are in novel shapes. For each object, we perform six independent trials with different ways of holding. The baseline chosen for comparison is GraspNet without depth restoration. In addition to assessing the overall success rate, we also compare the number of objects that exceed the specific success rate threshold $\delta \in \{0.5, 0.8, 1.0\}$.

TABLE III
RESULTS OF HANDOVER EXPERIMENT.

Methods	Success Attempts	Success Rate	$\delta_{0.5} \uparrow$	$\delta_{0.8} \uparrow$	$\delta_{1.0} \uparrow$
GraspNet	5/48	10.4%	0/8	0/8	0/8
HADR + GraspNet	34/48	70.8%	6/8	3/8	1/8

Table III reports the experimental results. Without depth restoration, GraspNet struggles to generate the correct grasp pose for handover due to the incomplete and inaccurate depth value. The proposed depth restoration method can significantly improve the handover success rate, from 10.4% to 70.8%. Six among eight objects achieve a success rate of over 50%.

V. CONCLUSION

In this study, we introduce a novel solution for restoring the depth of hand-held transparent objects and highlight its great value in human-robot interaction. We present TransHand-14K, a synthetic dataset generated through a real-to-sim data generation scheme, effectively bridging the gap between simulated and real-world environments. Our proposed method, HADR, leverages hand posture information to enhance depth restoration, achieving superior reconstruction results and better generalization capabilities compared to existing methods. Additionally, we develop a handover workflow that leverages our depth restoration technique. Evaluations on both depth restoration and human-to-robot handover validate the effectiveness of our approach.

REFERENCES

- [1] A. Mousavian, C. Eppner, and D. Fox, “6-DOF GraspNet: Variational grasp generation for object manipulation,” in *International Conference on Computer Vision (ICCV)*, 2019, pp. 2901–2910.
- [2] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, “3D diffusion policy,” *arXiv preprint arXiv:2403.03954*, 2024.
- [3] S. Sajjan, M. Moore, M. Pan, G. Nagaraja, J. Lee, A. Zeng, and S. Song, “ClearGrasp: 3D shape estimation of transparent objects for manipulation,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 3634–3642.
- [4] J. Jiang, G. Cao, J. Deng, T.-T. Do, and S. Luo, “Robotic perception of transparent objects: A review,” *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 6, pp. 2547–2567, 2023.
- [5] K. Chen, S. Wang, B. Xia, D. Li, Z. Kan, and B. Li, “TODE-Trans: Transparent object depth estimation with transformer,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 4880–4886.
- [6] Q. Dai, J. Zhang, Q. Li, T. Wu, H. Dong, Z. Liu, P. Tan, and H. Wang, “Domain randomization-enhanced depth simulation and restoration for perceiving and grasping specular and transparent objects,” in *European Conference on Computer Vision (ECCV)*, 2022, pp. 374–391.
- [7] W. Yang, C. Paxton, A. Mousavian, Y.-W. Chao, M. Cakmak, and D. Fox, “Reactive human-to-robot handovers of arbitrary objects,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 3118–3124.
- [8] S. Christen, W. Yang, C. Pérez-D’Arpino, O. Hilliges, D. Fox, and Y.-W. Chao, “Learning human-to-robot handovers from point clouds,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 9654–9664.
- [9] S. Albrecht and S. Marsland, “Seeing the unseen: Simple reconstruction of transparent objects from point cloud data,” in *Robotics: Science and Systems*, vol. 3, 2013.
- [10] C. J. Phillips, M. Lecce, and K. Daniilidis, “Seeing glassware: from edge detection to pose estimation and shape recovery,” in *Robotics: Science and Systems*, vol. 3. Michigan, USA, 2016, p. 3.
- [11] J. Ichnowski, Y. Avigal, J. Kerr, and K. Goldberg, “Dex-NeRF: Using a neural radiance field to grasp transparent objects,” *arXiv preprint arXiv:2110.14217*, 2021.
- [12] Q. Dai, Y. Zhu, Y. Geng, C. Ruan, J. Zhang, and H. Wang, “Grasp-NeRF: Multiview-based 6-dof grasp detection for transparent and specular objects using generalizable nerf,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1757–1763.
- [13] L. Zhu, A. Mousavian, Y. Xiang, H. Mazhar, J. van Eenbergen, S. Debnath, and D. Fox, “RGB-D local implicit function for depth completion of transparent objects,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 4649–4658.
- [14] H. Fang, H.-S. Fang, S. Xu, and C. Lu, “TransCG: A large-scale real-world dataset for transparent object depth completion and a grasping baseline,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7383–7390, 2022.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [16] NVIDIA, “Nvidia omniverse platform,” <https://developer.nvidia.com/nvidia-omniverse-platform>.
- [17] A. T. Miller and P. K. Allen, “Grasplit! a versatile simulator for robotic grasping,” *IEEE Robotics & Automation Magazine*, vol. 11, no. 4, pp. 110–122, 2004.
- [18] X. Weber, A. Xompero, and A. Cavallaro, “A mixed-reality dataset for category-level 6D pose and size estimation of hand-occluded containers,” *arXiv preprint arXiv:2211.10470*, 2022.
- [19] M. Meta, “Prime X Haptic - A haptic feedback glove,” 2024. [Online]. Available: <https://www.manus-meta.com/products/prime-x-haptic>
- [20] B. O. Community, “Blender - a 3D modelling and rendering package,” 2018. [Online]. Available: <https://www.blender.org/>
- [21] H. Yu, S. Li, H. Liu, C. Xia, W. Ding, and B. Liang, “TGF-Net: Sim2Real transparent object 6D pose estimation based on geometric fusion,” *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3868–3875, 2023.
- [22] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, “Swin Transformer: Hierarchical vision transformer using shifted windows,” in *International Conference on Computer Vision (ICCV)*, 2021, pp. 10012–10022.
- [23] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask R-CNN,” in *International Conference on Computer Vision (ICCV)*, 2017, pp. 2961–2969.
- [24] W. Yuan, T. Khot, D. Held, C. Mertz, and M. Hebert, “PCN: Point completion network,” in *International Conference on 3D vision (3DV)*, 2018, pp. 728–737.
- [25] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “PointNet: Deep learning on point sets for 3D classification and segmentation,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 652–660.
- [26] Y. Zhou and O. Tuzel, “VoxelNet: End-to-end learning for point cloud based 3D object detection,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4490–4499.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [28] W. Huang, P. Ren, J. Wang, Q. Qi, and H. Sun, “AWR: Adaptive weighting regression for 3D hand pose estimation,” in *The AAAI Conference on Artificial Intelligence*, vol. 34, no. 07, 2020, pp. 11 061–11 068.
- [29] C. Lugaresi, J. Tang, H. Nash, C. McClanahan, E. Uboweja, M. Hays, F. Zhang, C.-L. Chang, M. G. Yong, J. Lee *et al.*, “MediaPipe: A framework for building perception pipelines,” *arXiv preprint arXiv:1906.08172*, 2019.
- [30] H.-S. Fang, C. Wang, M. Gou, and C. Lu, “GraspNet-1Billion: A large-scale benchmark for general object grasping,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11 444–11 453.